Abstract: Let Σ denote a 2-dimensional surface. A graph G is irreducible for Σ provided that G does not embed into Σ, but every proper subgraph of G does. Let $I(\Sigma)$ denote the set of graphs with vertex degree at least three that are irreducible for Σ. In this paper we prove that $I(\Sigma)$ is finite for each orientable surface. Together with the result by D. Archdeacon and Huneke, stating that $I(\Sigma)$ is finite for each non-orientable surface, this settles a conjecture of Erdö's from the 1930s that $I(\Sigma)$ is finite for each surface Σ.

Let Σ_n denote the closed orientable surface of genus n. We also write $\gamma(\Sigma)$ to denote the genus of orientable surface Σ.

Let G be a finite graph. An embedding of G into a surface Σ is a topological map $\emptyset: G \to \Sigma$.

The orientable genus $\gamma(G)$ of the graph G is defined to be the least value of $\gamma(\Sigma)$ for all orientable surfaces Σ into which G can be embedded.

Let P be a property of a graph G. We say that G is P-critical provided that G has property P but no proper subgraph of G has property P. For example, if P is the property that $\gamma(G) \geq 1$, then the P-critical graphs are the two Kuratowski graphs K_5 and $K_{3,3}$. In general, if P is the property that $\gamma(G) \geq n$, then a P-critical graph can be embedded in Σ_n but not in Σ_{n-1}. Such a P-critical graph is also called irreducible for the surface Σ_{n-1}.

For any surface Σ, let $I(\Sigma)$ denote the set of graphs that have no vertices of degree two and are irreducible for Σ.

The result of this paper is the

Main Theorem: $I(\Sigma)$ is finite for each orientable surface Σ.

This result generalizes Kuratowski's theorem [9], that $I(\Sigma) = \{k_5, k_{3,3}\}$ if Σ is the sphere, to all orientable surfaces. The analogue result for non-orientable surfaces has been proved by D. Archdeacon and P. Huneke, who in [2] proved that $I_3(\Sigma)$, the set of all cubic graphs in $I(\Sigma)$, is finite for all non-orientable surfaces Σ and then extended in [1] their result to the full set $I(\Sigma)$ for all non-orientable surfaces Σ.

Together these results settle completely a conjecture of P. Erdös from the 1930's that $I(\Sigma)$ is finite for each closed two-dimensional surface Σ.

Some further historical references on this problem can be found in [2, 5-11]. There it is also noted that the finiteness of $I(\Sigma)$ is implied by the results of N. Robertson and P. Seymour on graph minors [12-14]. However, these results are far more general and difficult to derive than ours. The methods used in our proofs are more direct and in the spirit of Kuratowski's original proof.

In [1] the authors point out that their methods for proving the non-orientable analogue of our main theorem are not powerful enough to prove our main theorem. The reason is, loosely speaking, that the inductive step from Σ_n to Σ_{n+1} for orientable surfaces is twice as large -- measured in terms of the Euler characteristic -- as the step between Σ_{n+1} and Σ_n for non-orientable surfaces, where n denotes the orientable and non-orientable genus respectively. This larger step size markes the construction of all critical graphs for Σ_{n+1} from those for Σ_n more difficult in the orientable case.

We will show how this difficulty can be overcome by proving some new results, lemma 5 to lemma 11 while lemma 1 to lemma 4 can be obtained with little modification from the results in [1].

Before we give the proof of our main theorem, we mention a few definitions from [1] that we are going to use:

For a graph G with vertex set $V(G)$ and edge set $E(G)$, we call any vertex in $V(G)$ that is not of degree two in G a topological vertex of G and call any path of G whose endpoints are topological vertices of G while its innter vertices have all degree two in G a topological edge of G. A piece of G is either a topological vertex or the interior of a topological edge of G.

If H is a subgraph of G, then a bridge B of G with respect to H is the closure in G of a topological component of $G-H$. The vertices of attachment at B, denoted by v of $a(B)$, are the elements of $B \cap H$.

Proof of the main theorem

The proof of the main theorem is by induction on $\gamma(\Sigma)$, the orientable genus of Σ.

The base case, $\gamma(\Sigma) = 0$, is Kuratowski's theorem.

The induction step follows immediately from proposition 1 and the fact that each $(\gamma \geq n + 1)$-critical graph contains a $(\gamma \geq n)$-critical graph as subgraph. $\square$

Proposition 1: Let n be an arbitrary natural number and H be a $(\gamma \geq n)$-critical graph. Then there exist only finitely many $(\gamma \geq n + 1)$-critical graphs G that contain H as a subgraph.

Before giving the proof for this lemma, we want to point out that the proof is constructive in the sense that all the possible extensions of a $(\gamma \geq n)$-critical graph to a $(\gamma \geq n + 1)$-critical graph are investigated. As it turns out, different extensions lead often to isomorphic copies of the same $(\gamma \geq n + 1)$-critical graph. Also, a $(\gamma \geq n + 1)$-critical graph can contain different $(\gamma \geq n)$-critical subgraphs. It therefore appears to be quite difficult to obtain a tight upper bound on the number of non-isomorphic $(\gamma \geq n + 1)$-critical graphs. We therefore restrict our arguments to proving the finiteness of the number of all possible extensions of a $(\gamma \geq n)$-critical graph to a $(\gamma => + n + 1)$-critical graph.

We now come to the proof of proposition 1. First we observe that, because of the additive properties of $\gamma(G)$ for separable graphs [4], we can assume w.l.o.g. that G is 2-connected.

Further, by reasoning as in [1, 2], we can assume that each 2-connected $(\gamma \geq n + 1)$-critical graph is the proper extension of a 2-connected $(\gamma \geq n)$-critical graph K or of a 2-connected graph H that is obtained from a separable $(\gamma \geq n$-critical graph K in one of finitely many ways. In any case, each 2-connected $(\gamma \geq n + 1)$-critical graph is the extension of one of finitely many 2-connected graphs H. Let H^n denote the set of all these graphs.

To simplify notation, we denote by I_H^{n+1} the set of all ($\gamma \geq n + 1$-critical graphs that contain the graph H as a minor. Clearly, to prove proposition 1, it suffices to show that $|I_H^{n+1}|$ is finite for each graph H in H. To prove this fact, we use the following six lemmas, which we state for any graph H in H^n and any pair of graphs (G, H), where G belongs to I_H^{n+1}.

Lemma 1: Each embedding ϕ: $H \rightarrow \Sigma_n$ is an open 2-cell embedding.

Lemma 2: There exists an upper bound on the number of all (G, H) bridges that are attached to only one piece of G.

Lemma 3: There exists an upper bound on the number of all (G, H) bridges that are attached to at least three pieces of G.

Lemma 4: For each (G, H) bridge B, the number of topological vertices of B is bounded by $3 * \; |v\,\mathrm{of}a(B)|$.

These four lemmas follow readily from the results in [1, 2, 3, 15, 16].

As in the non-orientable case, it remains to prove the following two lemmas:

Lemma 5: There exists an upper bound on the number of all (G, H) bridges that are attached to two pieces of G.

Lemma 6: There exists an upper bound on $|v\,\mathrm{of}a(B)|$ for each (G, H) bridge B.

Clearly, proposition 1 follows directly from these six lemmas.

The proofs for lemmas 5 and 6 are not just generalizations of the analogue lemmas in the non-orientable case. They require some new concepts, which we will now outline.

We first observe that in order to prove lemma 6, it suffices, because of the finiteness of $|H^n|$, to show that there exists an upper bound on $|v\,\mathrm{of}a(B) \cap P|$ for any (G, H) bridge B and any piece P

of H. The proof for this fact can be further reduced to the proof of lemma 5.

The main burden in proving lemma 5 is the case that the two pieces of H are both topological edges of H. The other cases can be reduced to this one.

The proof of lemma 5 in that case is by contradiction. We will derive a contradiction in five steps, that are described by the following five lemmas. We first make the

Definition: Let e_1 and e_2 be two topological edges of a graph H in $\mathbf{H}^n$ and let G be a graph in $\mathbb{I}^{n+1}$. A set $\mathbf{B}$ of (G, H) bridges that are attached only to inner vertices of e_1 and e_2 is called a parallel bundle between e_1 and e_2 if there exists an embedding $\phi: H \rightarrow \Sigma_n$ for which all bridges in $\mathbf{B}$ can be embedded into one face of ϕ, leaving e_1 on one side only and arriving at e_2 from one side only. $|\mathbf{B}|$ is the width of the parallel bundle.

Remark 1: For such a parallel bundle $\mathbf{B}$ we choose an orientation of e_1 and number the bridges B_1 of $\mathbf{B}$ in such a way that $\mathbf{B}_1$ is always left of B_{i+1} for $1 \leq i \leq |\mathbf{B}| - 1$. Clearly, the orientation on e_1 induces a corresponding orientation on e_2.

Remark 2: If B_i is an inner bridge of $\mathbf{B}$, i.e., $2 \leq i \leq |\mathbf{B}| - 2$, and $\phi: G - B_i \rightarrow \Sigma_n$ is an embedding, then the bridges of $\mathbf{B}$ before and after B_i leave e_1 from opposite sides and arrive at e_2 from opposite sides, otherwise ϕ could be extended to an embedding of G into Σ_n, contradicting the criticality of G.

From the pigeon hole principle follows

Lemma 7: If there is a graph H in $\mathbf{H}^n$ and two topological edges e_1 and e_2 in H with the property that there exist graphs G in $\mathbb{I}^{n+1}$ such that the number of (G, H) bridges attached to e_1 and e_2 is unbounded, then there exist also graphs G in $\mathbb{I}^{n+1}$ with arbitrarily wide parallel bundles between e_1 and e_2.

This lemma can be extended to

Lemma 8: If there is a graph H in $\mathbf{H}$ and two topological edges e_1 and e_2 of H with the property that there exist graphs in $I\!\!H^{+1}$ with arbitrarily wide parallel bundles between e_1 and e_2, then there exists a topological edge e_3 of H with the property that there are graphs in $I\!\!H^{+1}$ with arbitrarily wide parallel bundles between e_1 and e_2 and between e_2 and e_3. Furthermore, the two parallel bundles are interlaced at e_2.

Continuing inductively, we make the following

Definition: Let H be a graph in $\mathbf{H}^n$ and $e_1, e_2, \cdots e_{l+1}$ topological edges of H, and let G be a graph in $I\!\!H^{+1}$. A sequence of parallel bundles $\mathbf{B}_i$ of G between e_i and e_{i+1}, $1 \le i \le l$, which are interlaced at e_i for $2 \le i \le l$, is called a string **S** of interlaced parallel bundles of G of length l.

Remark: For such a string **S**, the inner edges e_i, $2 \le i \le l$, must be pairwise distinct. If $e_1 = e_{l+1}$, then **S** is called a cycle **C** of interlaced parallel bundles of G of length l.

By induction, lemma 8 can be generalized to

Lemma 9: If there is a graph H in $\mathbf{H}^n$ and two topological edges e_1 and e_2 of H with the property that there exist graphs in $I\!\!H^{+1}$ with arbitrarily wide strings **S** of interlaced parallel bundles of length l between e_1 and e_2, then there exists a topological edge e_3 of H with the property that there are graphs in $I\!\!H^{+1}$ with arbitrarily wide string of interlaced parallel bundles of length $l + 1$ between e_1 and e_3. Omitting some details, lemma 9 says that wide strings of interlaced parallel bundles can be extended to longer strings.

Because of the finiteness of $|\mathbf{H}^n|$, the number of topological edges for any graph H in $\mathbf{H}^n$ is bounded. Therefore the length of strings of interlaced parallel bundles for any graph in $I\!\!H^{+1}$ is uniformly bounded. That means that the extension process described in lemma 9 has to stop after a bounded number of steps. The only way that this can happen is that this extension process leads eventually to a cycle of interlaced parallel bundles.

Combining the results of lemmas 7-9, we obtain

Lemma 10: From the assumptions of lemma 7 follows that there exist graphs in $I\!I^{n+1}$ with arbitrarily wide cycles cycles of interlaced parallel bundles.

To arrive at the contradiction needed to prove lemma 5, we now state

Lemma 11 There exists a uniform bound on the width of interlaced parallel bundles for all graphs in $I\!I^{n+1}$.

Proof: Let G be a graph in $I\!I^{n+1}$ and $\mathbf{C}$ a cycle of interlaced parallel bundles in G. If $\mathbf{C}$ is wide enough, we can derive a contradiction and so prove lemma 11.

We will show first that for any inner bridge B_i of a parallel bundle of $\mathbf{C}$ and any embedding $\phi: G - B_i \to \Sigma_n$, $\mathbf{C} - B_i$ must be embedded into two opposite cylinders, such that B_i cannot be embedded into either of them. This implies that the string of parallel bundles that forms $\mathbf{C}$ must be closed to a cycle in a unique way, forcing a unique interlacing of the first and last parallel bundle of this string.

We then conclude that G contains a subgraph H' isomorphic to H and a cycle $\mathbf{C}'$ of parallel bundles that is a proper sub-bundle of $\mathbf{C}$, closed in the same way as $\mathbf{C}$. For any edge e of $\mathbf{C} - \mathbf{C}'$, there exists an embedding $\phi: G - e \to \Sigma_n$. In this embedding, $\mathbf{C}'$ must be embedded into a single cyclinder. Clearly, by extending this embedding of $\mathbf{C}'$ to $\mathbf{C}$, the embedding of $G - e$ into Σ_n can be extended to an embedding of G into Σ_n, contradicting the criticality of G.

Details of this proof are best demonstrated on a cycle of length four. The case for general values of l follows by induction.

Let $\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3, \mathbf{B}_4$ be four parallel bundles of a graph G in $I\!I^{n+1}$, interlaced at the four topological edges e_2, e_3, e_4, e_1 of H respectively. W.l.o.g., assume that these edges are oriented in the same way and that the bridges in the four parallel bundles are numbered according with that orientation. Let B_{1i} be a bridge of $\mathbf{B}_1$ and let $\mathbf{B}_1^+$ and $\mathbf{B}_1^-$ denote the sub-bundles of $\mathbf{B}_1$ to the right and to the left of B_{1i} respectively, and assume w.l.o.g., that both contain at least two bridges.

Let ϕ: $G - B_{1i} \rightarrow \Sigma_n$ be an embedding. Clearly, $\mathbf{B}_1^+$ and $\mathbf{B}_1^-$ leave e_1 at opposite sides, w.l.o.g., assume B_1^+ leaves e_1 on the right (Figure 1). For B_{1i} not to be embeddable on the right side of e_1 parallel to B_1^+, there has to exist w.l.o.g. a first bridge B_{2i} of $\mathbf{B}_2$ with the property that B_{1i} and B_{2i} are interlaced at e_2 in such a way that the lowest vertex on e_2 is lower than the highest vertex of B_{2i} on e_2 and such that B_{2i} is embedded to the left of e_2 together with $\mathbf{B}_2^-$, consisting of all bridges of B_2 to the left of B_{2i}, while the remainder $\mathbf{B}_2^+$ of $\mathbf{B}_2$ is embedded to the right of e_2. Continuing to reason this way, we conclude that there has to exist a first bridge B_{3i} of $\mathbf{B}_3$, analogously to B_{2i}, blocking an embedding of B_{2i} next to $\mathbf{B}_2^+$ and again analogously, that there has to exist a first bridge B_{4i} of $\mathbf{B}_4$, blocking an embedding of B_{3i} next to $\mathbf{B}_3^+$. Again, B_{4i} leaves e_4 on the left together with $\mathbf{B}_4^-$, while $\mathbf{B}_4^+$ leaves e_4 on the right. In order that B_{4i} cannot be embedded parallel to $\mathbf{B}_4^+$ on the right of e_4 and on the left of e_1, the lowest vertex of B_{4i} on e_1 has to be lower than the highest vertex of the bridge of B_{1i} the bridges $B_{1\,i-1}$ in $\mathbf{B}_1^-$.

On the other hand, in order to block an embedding of B_{1i} on the left of e_1 parallel to $\mathbf{B}_1^-$, $B_{4\,i+1}$ must interfere; that means, the lowest vertex of $B_{4\,i+1}$ on e_1 has to be lower than the highest vertex of B_{1i} on e_1.

This argument can be applied to any inner bridge B_{1i} of $\mathbf{B}_1$ and shows that there is only one way that the parallel bundles $\mathbf{B}_1$ and $\mathbf{B}_4$ can be interlaced at e_1.

From these discussions it is obvious that a cycle of parallel bundles that is interlaced in this way cannot be embedded into two opposite cyclinders without excluding the embedding of one of its inner bridges.

We will show next how to construct a subgraph H' of G isomorphic to H and a proper sub-bundle $\mathbf{C}'$ of $\mathbf{C}$ that is connected to H' in the same way as $\mathbf{C}$ is connected to H.

For each bridge B_{ki}, $1 \le k \le 4$, let b_{ki} be a path in B_{ki} from the highest vertex of B_{ki} on e_k to the lowest vertex of B_{ki} on e_{k+1} (k mod 4).

We now define the new topological edges e_k' by

$$e_1' = e_{1l}\ b_{41}\ b_{32}\ b_{23}\ b_{14}\ e_{1r}$$

$$e_4' = e_{4l}\ b_{31}\ b_{22}\ b_{13}\ b_{45}\ e_{4r}$$

$$e_3' = e_{3l}\ b_{21}\ b_{12}\ b_{44}\ b_{35}\ e_{3r}$$

$$e_2' = e_{2l}\ b_{11}\ b_{43}\ b_{34}\ b_{25}\ e_{2r}.$$

Here e_{kl} and e_{kr} are the subpaths of e_k to the left and to the right of the paths b_{ji} respectively and the paths b_{kj} and $b_{k+1}\, j'$ are joined by appropriate subpaths of e_k (k mod 4), see Figure 2.

For $k = 2, 3, 4$ we define $\mathbf{B}_k'$ to be $\mathbf{B}_k^- \cup \mathbf{B}_k^{++}$, where $\mathbf{B}_k^-$ consists of all bridges of $\mathbf{B}_k$ to the left of B_{k1} and $\mathbf{B}^{k++}$ of all bridges of $\mathbf{B}_k$ to the right of $\mathbf{B}_{k5}$. Finally $\mathbf{B}_1'$ is defined as $\mathbf{B}_1^- \cup \mathbf{B}_1^{++} \cup \{p\}$, where $\mathbf{B}_1^-$ consists of all bridges of $\mathbf{B}_1$ to the left of B_{10}, $\mathbf{B}_1^{++}$ of all bridges to the right of B_{15} and p is the path $b_{10}\, b_{42}\, b_{33}\, b_{24}\, b_{15}$, where these paths are connected by appropriate subpaths of the topological edges e_k. (Figure 2)

It is straight forward to check that the four bundles $\mathbf{B}_k'$ form a cycle $\mathbf{C}'$ of parallel bundles that are interlaced in the same way as $\mathbf{C}$, completing the proof of lemma 11.

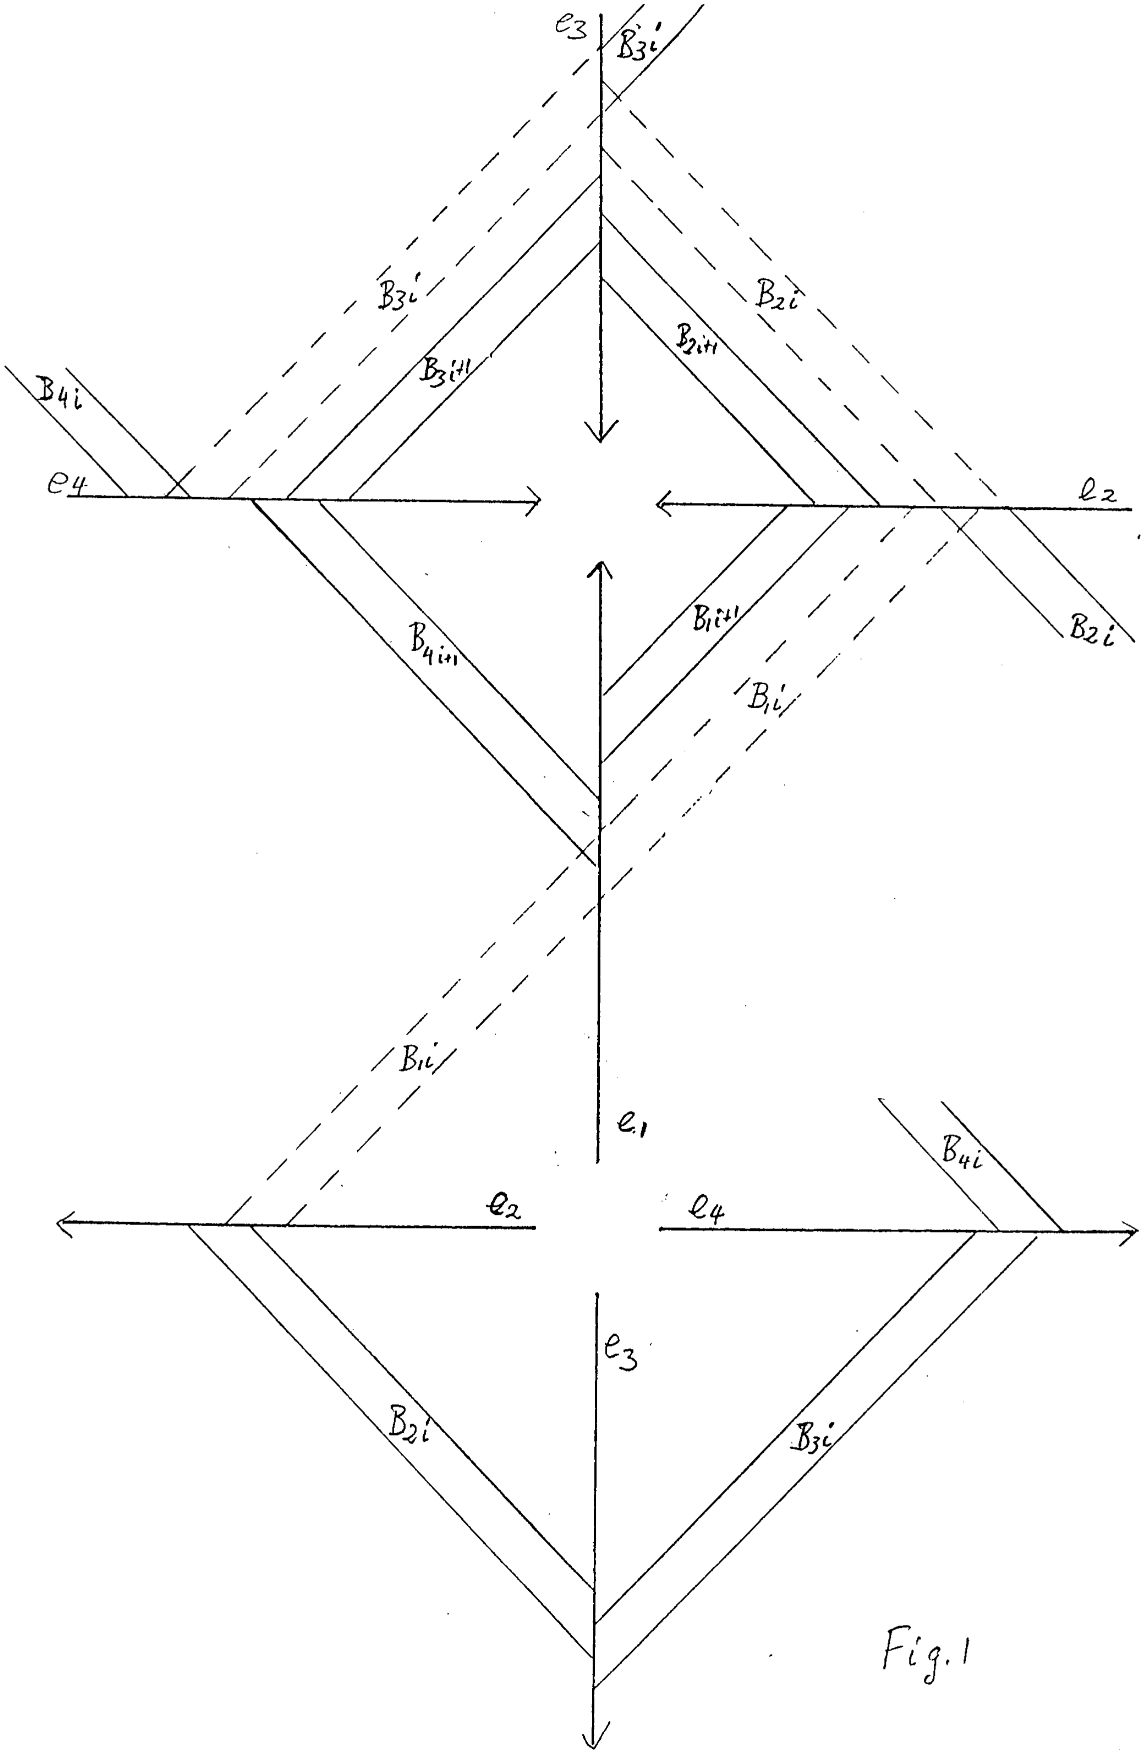
e3
B3i'
B3i'
B3i+1
B2i
B2i+1
B4i
e4
e2
B2i
B1i+1
B4i+1
B1i'
B1i'
e1
e2
e4
B4i
e3
B2i
B3i

Fig. 1

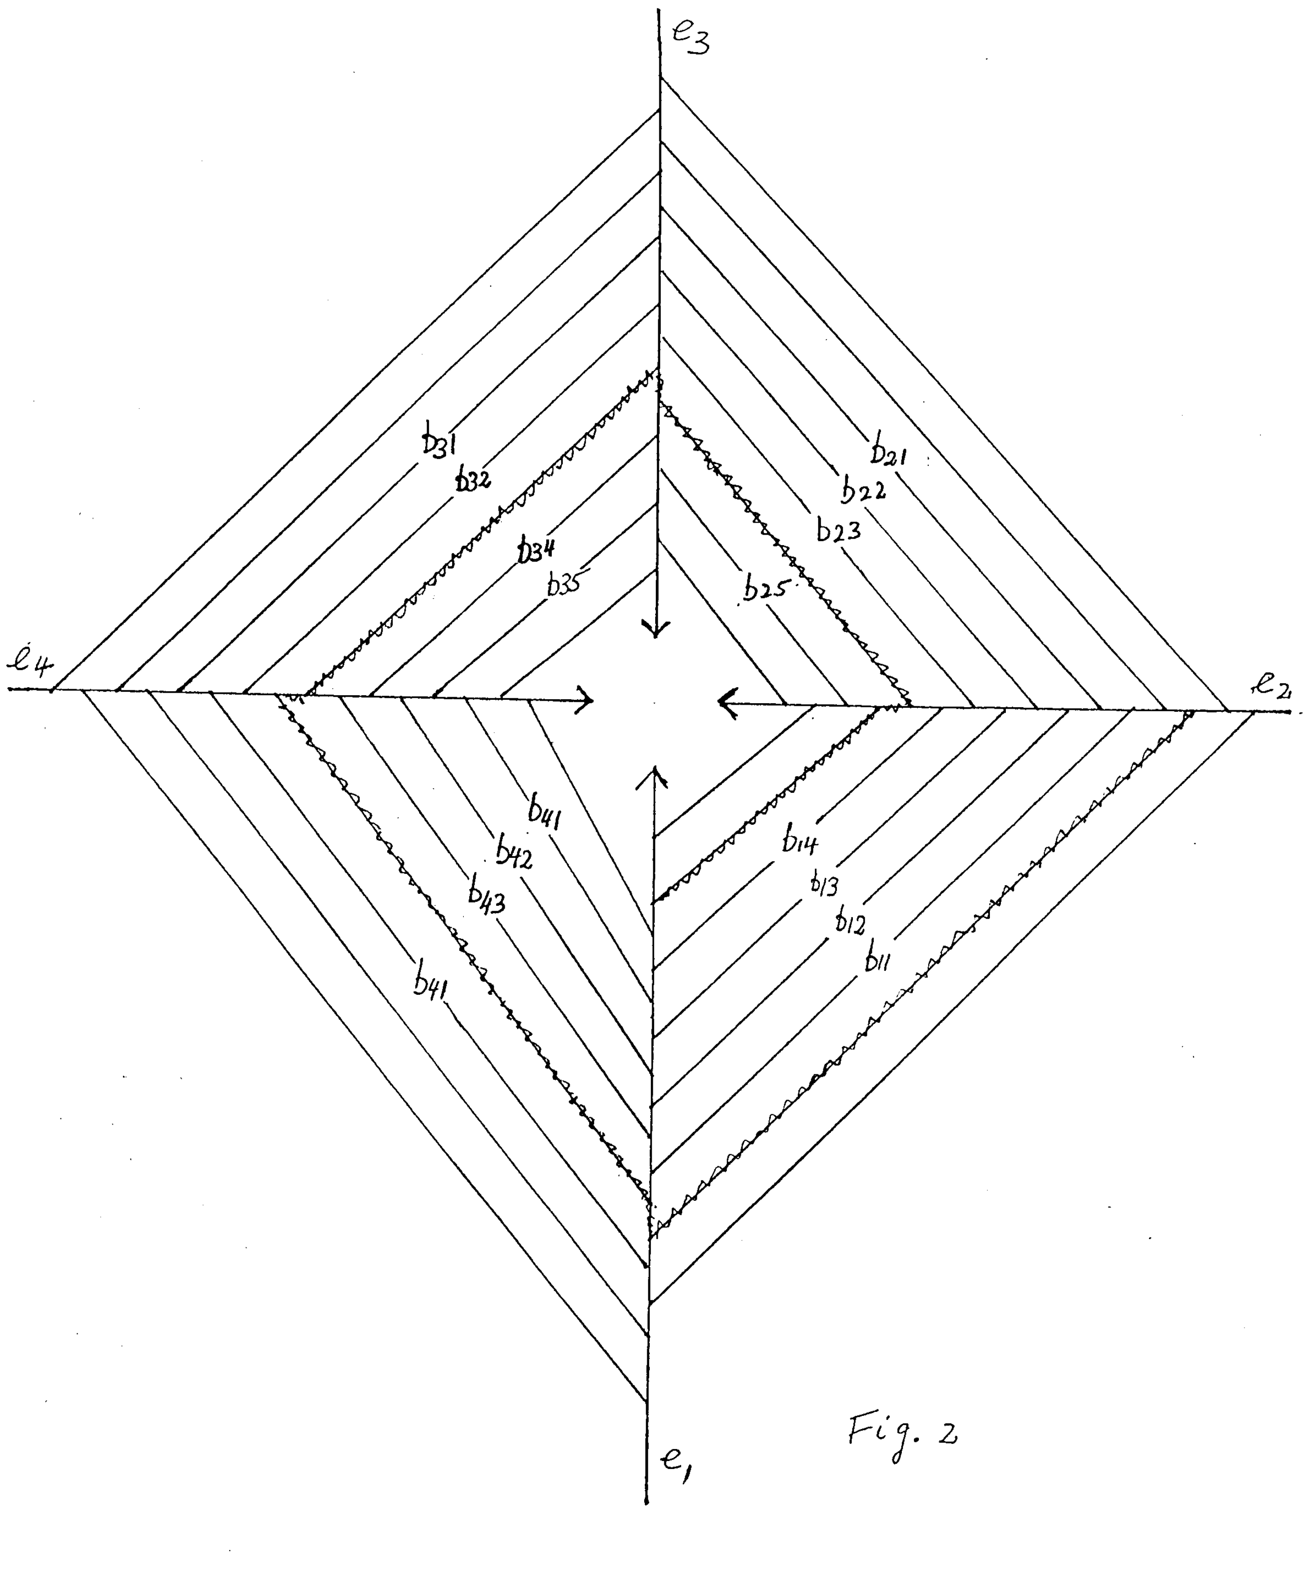

e3
e4
e2
e1
b31
b32
b34
b35
b21
b22
b23
b25
b41
b42
b43
b41
b14
b13
b12
b11

Fig. 2

.sh References[1-14]

1. D. Archdeacon and F. Huneke, *A Kuratowski Theorem for Nonorientable Surfaces,* preprint 1986.
2. D. Archdeacon and J. Huneke, "On Cubic Graphs which are Irreducible for Nonorientable Surfaces ," *J. Comb. Th. Ser. B.* **39**(3) pp. 233-264 (1985).
3. L. Auslander, T. Brown, and J.W.T. Youngs, "The Imbedding of Graphs in Manifolds," *J. of Math. and Mech.* **12**(4) pp. 229-234 ().
4. J. Battle, F. Harary, Y. Kodama, and J.W.T. Youngs, "Additivity of the Genus of a Graph," *Bull. Amer. Math. Soc.* **68** pp. 565-568 (1962).
5. H. Glover and J. Huneke, "Cubic Irreducible Graphs for the Projective Plane," *Discrete Math.* **13** pp. 341-355 (1975).
6. H. Glover and J. Huneke, "Graphs with Bounded Valence that do not Embed in Projective Plane," *Discrete Math.* **18** pp. 155-165 (1977).
7. H. Glover and J. Huneke, "The Set of Irreducible Graphs for the Projective Plane is Finite," *Discrete Math.* **22** pp. 243-256 (1978).
8. H. Glover, J. Huneke, and C.S. Wang, "103 Graphs that are Irreducible for the Projective Plane," *J. Comb. Th. Ser. B* **27** pp. 332-370 (1979).
9. K. Kuratowski, "Sur le probleme des Courbes Gauches en Topologie," *Fund. Math.* **15** pp. 271-283 (1930).
10. M. Miligram, "Irreducible Graphs," *J. Comb. Th. Ser. B.* **12** pp. 6-31 (1972).
11. M. Miligram, "Irreducible Graphs - Part 2," *J. Comb. Th. Ser. B.* **14** pp. 7-45 (1973).
12. N. Robertson and P. Seymour, *Graph Minors VIII: A Kuratowski Theorem for General Surfaces,* preprint
13. N. Robertson and P. Seymour, "Generalizing Kuratowski's Theorem," *Congressus Numerantium* **45** pp. 129-138 (1984).
14. N. Robertson and P. Seymour, "Graph Minors - A Survey," Surveys in Combinatorics, London Math. Soc. Lecture Note Ser. 105, Cambridge University Press (1975).